

REMARKS

This amendment is being filed along with a Request for Continued Examination (RCE) in response to the final Office Action having a mailing date of April 11, 2008. Various claims are amended as shown. No new matter has been added. With this amendment, claims 34-62 are pending in the application.

I. Supplemental information disclosure statement (IDS)

A supplemental IDS is being filed herewith to submit additional references. Because this supplemental IDS is being filed along with the present RCE, an IDS fee and/or an IDS certification are not needed and are therefore not being submitted. It is kindly requested that an Examiner-initialed copy of this supplemental IDS be provided along with the next communication, so as to confirm that the references listed therein have been entered into the record and considered.

II. Rejections under 35 U.S.C. § 112

The final Office Action rejected claim 62 under 35 U.S.C. § 112 for being indefinite. Claim 62 is amended as shown to address this indefiniteness rejection. In view of this amendment, it is kindly requested that this indefiniteness rejection be withdrawn.

III. Discussion of the claims and cited references

The final Office Action has rejected claims 34-62 under 35 U.S.C. § 102(b) as being anticipated by a white paper entitled “Server Load Balancing in Today’s Web-Enabled Enterprise” (hereinafter “the White Paper”).

The final Office Action also rejected claims 34-62 under 35 U.S.C. § 103(a) as being unpatentable over an Alteon document entitled “Enhancing Web User Experience with Global Server Load Balancing” (hereinafter “the Alteon document”) in view of the Cisco document entitled “Configuring the CSS Domain Name Server” (hereinafter “the Cisco document”).

For the reasons set forth below, these rejections are respectfully traversed.

A. Independent claim 34

Independent claim 34 as amended herein recites, *inter alia*, “obtaining at one of said site switches mapping information that provides a translation between a private virtual IP address, configured at said site switch and associated with said at least one host server corresponding to said site switch, and a public virtual IP address” and “providing, by said site switch, said public virtual IP address to at least one load balancing controller.” It is respectfully submitted that the final Office Action has not cited any passage from the references that meets these specific limitations.

For example, page 3 (section 2) of the final Office Action stated the following in discussing the White Paper (emphasis ours):

“When a local DNS looks up an authorized DNS via a controller GSLB switch, the response is returned to the client with the IP address information from the controller GSLB switch via the local DNS. This does the act: ‘providing said mapping public virtual IP address from said site switch’ because it provides the client to establish a connection within the site in London or a site in Hong Kong.”

From the above-quoted passage from the final Office Action, the Examiner has clearly interpreted the “controller GSLB switch” as performing the act of “providing … said public virtual IP address…” However, claim 1 as amended herein clarifies that it is the site switch that performs the recited “providing said public virtual IP address to at least one load balancing controller.” Accordingly, since the White Paper does not meet the limitations of claim 1 that require “providing, by said site switch, said public virtual IP address to at least one load balancing controller,” claim 1 is allowable over the White Paper.

Claim 1 also recites the limitation of “obtaining at one of said site switches mapping information that provides a translation between a private virtual IP address, configured at said site switch and associated with said at least one host server corresponding to said site switch, and a public virtual IP address.” It is respectfully submitted that this limitation is also not met by the White Paper.

For example, page 5 (section 6) of the final Office Action has cited the Figure on Page 6 of the White Paper as allegedly meeting this limitation of claim 1. More particularly, the final Office Action cited “the Figure – a Site such as a box 4 in the Figure, having a Controller GSLB switch that configures to associate with at least one host server: Hong Kong.” However, the “box 4” and the accompanying description of the White Paper being relied upon by the final Office Action is completely silent with respect to “obtaining at one of said site switches mapping information that provides a translation between a private virtual IP address, configured at said site switch and associated with said at least one host server corresponding to said site switch, and a public virtual IP address”

It is respectfully submitted that with the technology of the White Paper, for a situation where a private virtual IP address was configured at the site switch, that site switch did not obtain the mapping information that provides a translation between the private virtual IP address and the public virtual IP address.

As further evidence of the allowability of claim 34 over the subject matter disclosed in the White Paper, included herewith along with this present amendment/response is a September 30, 2008 affidavit by Prajakta S. Joshi, the contents of which are incorporated herein by reference. As stated in the affidavit by Ms. Joshi, who is knowledgeable of the subject matter described in the White Paper, for an architecture where a private virtual IP address was configured at the site switch, the site switch did not communicate public virtual IP addresses to the controller GSLB switch (CGS) of the White Paper, prior to her invention. Instead for an architecture where a private virtual IP address was configured at the site switch, the private virtual IP address (rather than a public virtual IP address) was communicated by the site switch to the CGS of the White Paper. With the embodiments of her invention, the site switch obtained

(from a mapping device such as a NAT device or a firewall device) the mapping information that provides a translation between the private virtual IP address and the public virtual IP address, thereby enabling the site switch to provide this public virtual IP address to the load balancing controller.

Accordingly, it is respectfully submitted that claim 34 is thus further allowable over the White Paper.

Claim 34 is also allowable over the Alteon document and the Cisco document. Ms. Joshi's affidavit provides further information/proof as to how the embodiments of her invention are different from the subject matter of the Alteon document and the Cisco document.

For example, Figure One and the accompanying description on page 2 of the Alteon document describe site switch A that "returns site B's virtual IP address (VIP) address [172.176.110.20] to the client's local DNS." The local DNS server then "responds to client with site B's VIP" and the client "opens application session to IP 172.176.110.20." Since the VIP address 172.176.110.20 is returned to the client and the client is able to open a session to this VIP address, this means that the VIP address 172.176.110.20 is a public VIP address configured at site B (described in the Alteon document as "site B's virtual IP address"). Thus, the Alteon document does not describe an implementation involving private VIP address configured at the site switch, and therefore it is respectfully submitted that the Alteon document does not provide the features in claim 1 of (a) "obtaining at one of said site switches mapping information that provides a translation between a private virtual IP address, configured at said site switch and associated with said at least one host server corresponding to said site switch, and a public virtual IP address," and (b) "providing, by said site switch, said public virtual IP address to at least one load balancing controller."

Pages 12-13 of the Cisco document describe the configuration to use a content services switch (CSS) to perform network address translation (NAT) to translate a private IP address of a server (*e.g.*, the private IP address 10.0.3.251) to a public VIP address (*e.g.*, the public VIP address 192.200.200.200) and vice versa. According to Ms. Joshi's reading of the Cisco document, the private IP address 10.0.3.251 of the server described in the Cisco document

is a private real IP address of the server, rather than a private virtual IP address that is configured at a site switch. Evidence that the private IP address 10.0.3.251 is a real IP address of a server, rather than a private VIP address configured at a site switch, is provided on page 12 of the Cisco document, which states that “The source group enables the CSS to perform Network Address Translation to translate outbound traffic source IP addresses from the server’s private IP address (10.0.3.251) to the public VIP address (192.200.200.200). To prevent server source port collisions, the CSS performs Network Address Translation on the server’s source IP address and port by translating the: Source IP address to the IP address defined in the source group” (emphasis added). Accordingly, the Cisco document also does not provide the claimed features in claim 1 of (a) “obtaining at one of said site switches mapping information that provides a translation between a private virtual IP address, configured at said site switch and associated with said at least one host server corresponding to said site switch, and a public virtual IP address,” and (b) “providing, by said site switch, said public virtual IP address to at least one load balancing controller.”

Hence, claim 1 is further allowable over the cited references.

B. Independent claims 39, 43, 47, 51, and 56

Independent claims 43 and 51 recite limitations (using varying language) generally corresponding to those of claim 34, for a site switch. Independent claims 39, 47, and 56 are written from the point of view of a load balance switch, and recite for example, receiving a public virtual IP address and an address record that indicates the received public virtual IP address as being associated with the site switch.

For reasons similar to those explained above, it is respectfully submitted that claims 39, 43, 47, 51, and 56 are allowable as well.

C. Dependent claims 37, 40, 48, 54, and 57

Dependent claim 37 as amended herein recites, *inter alia*, that virtual public IP addresses received by the load balancing controller that “do not have indication in an address

record as being associated with corresponding said site switches ... are excluded from having applied thereto any metric of a load balancing algorithm that is usable with virtual IP addresses."

It is respectfully submitted that the final Office Action has failed to cite any passage of the White Paper and the other references that meets these limitations.

For example in rejecting previously presented claim 37 on the basis of the White Paper, page 6 of the final Office Action stated "See the operation in the Figure in [page] 6, and look up process. Particularly, in the paragraph in p. 5 'SLB Technique', discussing about allowing the client to access as a real server for subsequent requests." This reliance on the White Paper to reject claim 37 is traversed herein.

For example, nothing in the cited passages of the White Paper disclose, teach, or suggest the exclusion of application of a metric of load balancing algorithm (usable with virtual IP addresses), to public virtual IP addresses that do not have indication in an address record as being associated with the site switch(es). If the Examiner believes that such teaching is indeed present in the White Paper, then the Examiner is kindly requested to specifically identify/quote such teachings from the White Paper, since a review of the White Paper by the undersigned attorney has failed to locate such teachings that involve excluding the application of a metric to public VIP addresses that do not have indication in an address record as being associated with a site switch.

Hence, claim 37 is allowable over the White Paper.

In rejecting claim 37, page 12 of the final Office Action has also cited page 2 of the Alteon document. Specifically, the final Office Action relies upon the Alteon document, by stating "See p. 2, including three bullets, GSLB, develops a list or order list of sites, including site health, and geographic locations, and see all last three paragraphs." However, it is respectfully submitted herein that the Alteon document, contrary to the position taken by the final Office Action, does not meet the limitations of claim 37.

First, the criteria of "site health," "geographic location," and "response time" taught by the Alteon document and relied upon by the final Office Action are merely metrics, and teach absolutely nothing regarding "public virtual IP addresses ... that do not have indication

in an address record as being associated with corresponding said site switches ... are excluded from having applied thereto any metric of a load balancing algorithm that is usable with virtual IP addresses.” Stated in another way, the cited passages of Alteon relating to these metrics do not teach the exclusion of claim 37.

Second, the last three paragraphs of the Alteon document teach that there is a “redirect” to a different site if “the site to which the client has been pointed suddenly experiences a failure or is overloaded” (emphasis ours). Since the client “has been pointed” to the site, then the address clearly is associated with that site and a metric of a load balancing algorithm necessarily has been applied to the address of that site, which has now failed or become overloaded. This passage of the Alteon document therefore does not meet the limitations of claim 37 that require “public virtual IP addresses ... that do not have indication in an address record as being associated with corresponding said site switches ... are excluded from having applied thereto any metric of a load balancing algorithm that is usable with virtual IP addresses.”

Hence, claim 37 is further allowable.

Dependent claims 40, 48, 54, and 57 recite limitations (using varying language) generally corresponding to claim 37, and are allowable for similar reasons.

D. Dependent claims 38, 42, 46, 50, 55, and 58

Dependent claim 38 as amended herein recites, *inter alia*, “an active bindings metric that prefers a virtual IP address, configured at respective said site switches, having a maximum number of active ones of said host servers bound to said preferred virtual IP address, rather than preferring another virtual IP address having a number of bound active ones of said host servers that is less than said maximum number.” In rejecting previously presented claim 38, page 6 of the final Office Action continues to cite page 9 (the “High Availability” and “Maximum Scalability” subsections) of the White Paper. It is respectfully submitted that the final Office Action has not adequately identified any teachings in these cited passages of the White Paper that meet the specific limitations for an active bindings metric as recited in claim 38.

The passages on page 9 of the White Paper relied upon by the final Office Action are reproduced in full below:

“High Availability for Server and Application

ServerIron switches ensure service availability by offering switch, server, link, and session level redundancy. In the event of a server or application outage, ServerIron switches provide detection and sub-second fail-over to the next server that supports a like service. ServerIron switches provide active-standby redundancy or active-active redundancy capability that allows administrators to establish primary and secondary load balancing switches to support identical configurations parameters and provide 100 percent availability.

Maximum Scalability

TrafficWorks™ IronWare™ running on ServerIron simplifies network design by enabling IT managers to create a server farm, represented by a single IP address known as a virtual IP address (VIP). ServerIron acts as a single server for incoming Web traffic, controlling, monitoring and directing client requests to the most appropriate real server in a server farm. ServerIron’s firewall load balancing (FWLB) eliminates firewall bottlenecks by distributing load across multiple firewalls. FWLB is Check Point certified, and can load balance up to 32 firewalls for a scalable and highly available Web-site deployment.”

As clearly evident from the above-quoted passages from the White Paper, the “active bindings metric” as recited in claim 38 is not disclosed, taught, or suggested by these passages.

For example, the “High Availability” subsection is speaking of redundancy capability in which a “primary” load balancing switch is supported by a redundant “secondary” load balancing switch, in the event that the primary load balancing switch experiences a failure. It is respectfully submitted that the final Office Action has not explained how such load balancing switch “redundancy” teachings of the White Paper meet the limitations of claim 38 for an active bindings metric that prefers a virtual IP address (configured at a site switch) having a maximum number of active host servers bound to it, rather than another virtual IP address having a lesser number of active host servers. Stated in another way, the cited passage of the White Paper is speaking of providing redundancy between load balancing switches, whereas claim 38 is directed towards a load balancing controller that applies an active bindings metric that prefers a particular virtual IP address over another virtual IP address. The cited passage and claim 38 are thus addressing two different concepts.

As another example, the “Maximum Scalability” subsection is speaking of “firewall load balancing” (FWLB) in which traffic bottlenecks are eliminated by distributing the traffic across multiple firewalls. Again, it is respectfully submitted that the final Office Action has not explained how such FWLB teachings of the White Paper meet the limitations of claim 38 for an active bindings metric that prefers a virtual IP address (configured at a site switch) having a maximum number of active host servers bound to it, rather than another virtual IP address having a lesser number of active host servers. Stated in another way, even if the firewalls of the White Paper might be *hypothetically* associated with virtual IP addresses, the White Paper makes no mention of preferring one virtual IP address over another—the White Paper simply says that load is “distributed across multiple firewalls” and provides no further detail as to any metrics that may be used to decide how such load is distributed.

Accordingly, it is respectfully submitted that claim 38 is allowable.

Page 12 final Office Action also relies upon the Alteon document as allegedly meeting the limitations of claim 38. In particular, the final Office Action alleges that the discussion regarding the “maximum connection thresholds” in page 2 of the Alteon document

meets the active bindings metric limitations of claim 38. This allegation by the final Office Action is traversed herein.

More particularly, page 2 (last 3 paragraphs) of the Alteon document is speaking of “redirecting” the client to a different site if there is an overload at the site where the client was pointed. The Alteon document then explains that an overload occurs “when an HTTP request is sent to real servers that have reached their respective maximum connection thresholds.” Accordingly, the “maximum connection threshold” of Alteon relied upon by the final Office Action pertains to the amount of traffic that a server/site can carry, before an overload occurs. This is completely different from what is recited in claim 38, wherein the load balancing algorithm uses “an active bindings metric that prefers a virtual IP address, configured at respective said site switches, having a maximum number of active ones of said host servers bound to said preferred virtual IP address, rather than preferring another virtual IP address having a number of bound active ones of said host servers that is less than said maximum number.”

Hence, claim 38 is further allowable.

Claims 42, 46, 50, 55, and 58 recite limitations (using varying language) generally corresponding to claim 38, and are allowable for similar reasons.

E. Other claim amendments

Various other amendments are made to the claims as shown to provide appropriate antecedent basis, to make the language between and within the claims consistent given the amendments to the independent claims, to remove redundant and/or unnecessary limitations, to make changes of a typographical/grammatical nature, to more precisely recite the subject matter contained in the claims, and/or to otherwise place such claims in better form.

IV. Conclusion

If there are any informalities or questions that can be addressed via telephone, the Examiner is encouraged to contact the attorney of record (Dennis M. de Guzman) at (206) 622-4900.

The Director is authorized to charge any additional fees due by way of this Amendment, or credit any overpayment, to our Deposit Account No. 19-1090.

All of the claims remaining in the application are believed to be allowable. Favorable consideration and a Notice of Allowance are earnestly solicited.

Respectfully submitted,

SEED Intellectual Property Law Group PLLC

/Dennis M. de Guzman/

Dennis M. de Guzman
Registration No. 41,702

DMD:wt

701 Fifth Avenue, Suite 5400
Seattle, Washington 98104
Phone: (206) 622-4900
Fax: (206) 682-6031

1162708_1.DOC